

Timed Multiset Rewriting and the Verification of Time-Sensitive Distributed Systems

Max Kanovich^{1,5} Tajana Ban Kirigin² Vivek Nigam³ Andre Scedrov^{4,5} and Carolyn Talcott⁶

¹ Queen Mary, University of London & University College, UK, mik@dcs.qmul.ac.uk

² University of Rijeka, HR, bank@math.uniri.hr

³ Federal University of Paraba, João Pessoa, Brazil, vivek@ci.ufpb.br

⁴ University of Pennsylvania, Philadelphia, USA, scedrov@math.upenn.edu

⁵ National Research University Higher School of Economics, Moscow, Russia

⁶ SRI International, USA, clt@csl.sri.com

Abstract. Time-Sensitive Distributed Systems (TSDS), such as applications using autonomous drones, achieve goals under possible environment interference (*e.g.*, winds). Moreover, goals are often specified using explicit time constraints which must be satisfied by the system *perpetually*. For example, drones carrying out the surveillance of some area must always have *recent pictures*, *i.e.*, at most M time units old, of some strategic locations. This paper proposes a Multiset Rewriting language with explicit time for specifying and analysing TSDSes. We introduce two properties, *realizability* (some trace is good) and *survivability* (where, in addition, all admissible traces are good). A good trace is an infinite trace in which goals are perpetually satisfied. We propose a class of systems called *progressive timed systems* (PTS), where intuitively only a finite number of actions can be carried out in a bounded time period. We prove that for this class of systems both the realizability and the survivability problems are PSPACE-complete. Furthermore, if we impose a bound on time (as in bounded model-checking), we show that for PTS, realizability becomes NP-complete, while survivability is in the Δ_2^P class of the polynomial hierarchy. Finally, we demonstrate that the rewriting logic system Maude can be used to automate time bounded verification of PTS.

1 Introduction

The recent years have seen an increasing number of applications where computing is carried out in all sorts of environments. For example, drones are now being used to carry out tasks such as delivering packages, monitoring plantations and railways. While these distributed systems should still satisfy well-known safety (*e.g.*, drones should not run out of energy) and liveness properties (*e.g.*, freedom of livelock), they are also subject to *quantitative constraints* leading to new verification problems with explicit time constraints.

Consider, as our running example, the scenario where drones monitor some locations of interest such as infested plantation areas⁷, whether rail tracks are in place⁸, or locations

⁷ See <http://www.terradrone.pt/>) – In Portuguese.

⁸ See <http://fortune.com/2015/05/29/bnsf-drone-program/>.

with high risk of being trespassed. Drones should take a picture of each one of these points. Moreover, for each point, there should be a *recent picture*, i.e., not more than M time units old for some given M . That is, the drones should collectively have a set of *recent pictures* of all sensitive locations. In order to achieve this goal, drones may need to fly consuming energy and they may need to return to the base station to recharge their batteries. The environment may interfere as there might be winds that may move the drone to some direction or other flying objects that may block a drone's progression.

When designing such a system, engineers should specify the behavior of drones, e.g., where to move, when to take a picture, when to return to a base station, etc. A verification problem, called *realizability problem*, is to check, whether under the given time constraints, the specified system can achieve the assigned goal, e.g., always collect a recent picture of the sensitive locations.

In many settings, the drones themselves or the environment may behave non-deterministically. For example, if a drone wants to reach a point to the northeast, it may first choose to either move north or east, both being equally likely. Similarly, there might be some wind at some location causing any drone under the wind's effect to move in the direction of the wind. A stronger property that takes into account such non-determinism is to check whether for all possible outcomes (of drone actions or environment interference), the specified system can achieve the assigned goal. We call this property *survivability*.

In our previous work [13,15,12], we proposed a timed Multiset Rewriting (MSR) framework for specifying compliance properties which are similar to *quantitative safety properties* investigating the complexity of a number of decision problems. These properties were defined over the set of *finite traces*, i.e., the execution of a finite number of actions. Realizability and survivability, on the other hand, are similar to *quantitative liveness problems*, defined over infinite traces.

The transition to properties over infinite traces leads to many challenges as one can easily fall into undecidability fragments of verification problems. A main challenge is to identify the syntactical conditions on specifications so that the survivability and feasibility problems fall into a decidable fragment and at the same time interesting examples can be specified. Also the notion that a system satisfies a property perpetually implies that the desired property should be valid at all time instances independent of environment interference. Another issue is that systems should not be allowed to perform an unbounded number of actions in a single time instance a problem similar to the Zeno paradox.

The main contribution of this paper is threefold:

1. We propose a novel class of systems called *progressive timed systems* (PTS) (Section 2), specified as timed MSR theories, for which, intuitively, only a finite number of actions can be carried out in a bounded time. We demonstrate that our drone example belongs to this class (Section 3). We define a language for specifying realizability and survivability properties (Section 4) demonstrating that many interesting problems in Time-Sensitive Distributed Systems (TSDS) can be specified using our language;
2. We investigate (Section 5) the complexity of deciding whether a given system satisfies realizability and survivability. While these problems are undecidable in general, we show that they are PSPACE-complete for PTS. We also show that when we bound time

(as in bounded-model checking) the realizability problem for PTS is NP-complete and survivability is in the Δ_2^P class of the polynomial hierarchy (P^{NP}) [21].

3. Finally (Section 6), we show that the rewriting logic tool Maude [6] can be used to automate the analysis of TSDS. We implemented the drone scenario described above following the work of Talcott *et al.* [24] and carried out a number of simulations with different instances of this scenario. Our simulations demonstrate that specifiers can quickly find counter-examples where their specifications do not satisfy time bounded survivability.

We conclude by discussing related and future work (Section 7).

2 Timed Multiset Rewriting Systems

Assume a finite first-order typed alphabet, Σ , with variables, constants, function and predicate symbols. Terms and facts are constructed as usual (see [9]) by applying symbols of correct type (or sort). We assume that the alphabet contains the constant $z : Nat$ denoting zero and the function $s : Nat \rightarrow Nat$ denoting the successor function. Whenever it is clear from the context, we write n for $s^n(z)$ and $(n + m)$ for $s^n(s^m(z))$.

Timestamped facts are of the form $F@t$, where F is a fact and $t \in \mathbb{N}$ is natural number called *timestamp*. (Notice that timestamps are *not* constructed by using the successor function.) There is a special predicate symbol *Time* with arity zero, which will be used to represent global time. A *configuration* is a multiset of ground timestamped facts, $\mathcal{S} = \{Time@t, F_1@t_1, \dots, F_n@t_n\}$, with a single occurrence of a *Time* fact. Configurations are to be interpreted as states of the system. Consider the following configuration where the global time is 4.

$$\mathcal{S}_1 = \{Time@4, Dr(d1, 1, 2, 10)@4, Dr(d2, 5, 5, 8)@4, P(p1, 1, 1)@3, P(p2, 5, 6)@0\} \quad (1)$$

Fact $Dr(dId, x, y, e)@t$ denotes that drone dId is at position (x, y) at time t with e energy units left in its battery; fact $P(pID, x, y)@t$ denotes that the point to be monitored by pID is at position (x, y) and the last picture of it was taken at time t . Thus, the above configuration denotes a scenario with two drones at positions $(1, 2)$ and $(5, 5)$ and energy left of 10 and 8, and two points to be monitored at positions $(1, 1)$ and $(5, 6)$, where the former has been taken a photo at time 3 and the latter at time 0.

Configurations are modified by multiset rewrite rules which can be interpreted as actions of the system. There is only one rule that modifies global time:

$$Time@T \longrightarrow Time@(T + 1) \quad (2)$$

where T is a time variable. Applied to a configuration, $\{Time@t, F_1@t_1, \dots, F_n@t_n\}$, it advances global time by one, resulting in $\{Time@(t + 1), F_1@t_1, \dots, F_n@t_n\}$.

The remaining rules are *instantaneous* as they do not modify global time, but may modify the remaining facts of configurations (those different from *Time*). Instantaneous rules have the form:

$$Time@T, \mathcal{W}, F_1@T'_1, \dots, F_n@T'_n \mid \mathcal{C} \longrightarrow Time@T, \mathcal{W}, Q_1@(T + D_1), \dots, Q_m@(T + D_m) \quad (3)$$

where D_1, \dots, D_m are natural numbers, $\mathcal{W} = W_1@T_1, \dots, W_n@T_n$ is a set of times-tamped predicates possibly with variables, and \mathcal{C} is the guard of the action which is a set of constraints involving the variables appearing in the rule's pre-condition, *i.e.* the variables $T, T_1, \dots, T_p, T'_1, \dots, T'_n$. Following [8] we say that $F_1@T'_1, \dots, F_n@T'_n$ are consumed by the rule and $Q_1@(T + D_1), \dots, Q_m@(T + D_m)$ are created by the rule. (In a rule, we color **red** the consumed facts and **blue** the created facts.)

Constraints may be of the form:

$$T > T' \pm N \quad \text{and} \quad T = T' \pm N \quad (4)$$

where T and T' are time variables, and $N \in \mathbb{N}$ is a natural number. All variables in the guard of a rule are assumed to appear in the rule's pre-condition. We use $T \geq T' \pm N$ to denote the disjunction of $T > T' \pm N$ and $T = T' \pm N$.

A rule $W \mid \mathcal{C} \longrightarrow W'$ can be *applied on a configuration* \mathcal{S} if there is a ground substitution σ , such that $W\sigma \subseteq \mathcal{S}$ and $\mathcal{C}\sigma$ is true. The resulting configuration is $(\mathcal{S} \setminus W) \cup W'\sigma$. We write $\mathcal{S} \longrightarrow_r \mathcal{S}_1$ for the one-step relation where configuration \mathcal{S} is rewritten to \mathcal{S}_1 using an instance of rule r .

Definition 1. A *timed MSR system* \mathcal{A} is a set of rules containing only *instantaneous rules* (Equation 3) and the *tick rule* (Equation 2).

A *trace* of a timed MSR \mathcal{A} starting from an initial configuration \mathcal{S}_0 is a sequence of configurations where for all $i \geq 0$, $\mathcal{S}_i \longrightarrow_{r_i} \mathcal{S}_{i+1}$ for some $r_i \in \mathcal{A}$.

$$\mathcal{S}_0 \longrightarrow \mathcal{S}_1 \longrightarrow \mathcal{S}_2 \longrightarrow \dots \longrightarrow \mathcal{S}_n \longrightarrow \dots$$

In the remainder of this paper, we will consider a particular class of timed MSR, called *progressive timed MSR* (PTS), which are such that only a finite number of actions can be carried out in a bounded time interval which is a natural condition for many systems. We built PTS over balanced MSR taken from our previous work [16]. The balanced condition is necessary for decidability of problems (such as reachability as well as the problems introduced in Section 4).

Definition 2. A *timed MSR* \mathcal{A} is *balanced* if for all instantaneous rules $r \in \mathcal{A}$, r creates the same number of facts as it consumes, that is, in Eq. (3), $n = m$.

Proposition 1. Let \mathcal{A} be a balanced timed MSR. Let \mathcal{S}_0 be an initial configuration with exactly m facts. For all possibly infinite traces \mathcal{P} of \mathcal{A} starting with \mathcal{S}_0 , all configurations \mathcal{S}_i in \mathcal{P} have exactly m facts.

Definition 3. A *timed MSR* \mathcal{A} is *progressive* if \mathcal{A} is balanced and for all instantaneous rules $r \in \mathcal{A}$:

- rule r creates at least one fact with timestamp greater than the global time, that is, in Equation (3), at least one $D_i \geq 1$;
- rule r consumes only facts with timestamps in the past or at the current time, that is, in Equation (3), the set of constraints \mathcal{C} contains the set $\mathcal{C}_r = \{T \geq T'_i \mid F_i@T'_i, 1 \leq i \leq n\}$.

The following proposition establishes a bound on the number of instances of instantaneous rules appearing between two consecutive instances of Tick rules, while the second proposition formalizes the intuition that PTS always move forward.

$$\begin{aligned}
& \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y, E+1)@T \mid \text{doMove}(Id, X, Y, E+1, T, T_1, \dots, T_n, \text{north}) \longrightarrow \\
& \quad \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y+1, E)@(T+1) \\
& \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y+1, E+1)@T \mid \text{doMove}(Id, X, Y+1, E+1, T, T_1, \dots, T_n, \text{south}) \longrightarrow \\
& \quad \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y, E)@(T+1) \\
& \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X+1, Y, E+1)@T \mid \text{doMove}(Id, X+1, Y, E+1, T, T_1, \dots, T_n, \text{west}) \longrightarrow \\
& \quad \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y, E)@(T+1) \\
& \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y, E+1)@T \mid \text{doMove}(Id, X, Y, E+1, T, T_1, \dots, T_n, \text{east}) \longrightarrow \\
& \quad \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, X, Y, E)@(T+1) \\
& \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, x_b, y_b, E)@T \mid \text{doCharge}(Id, E, T, T_1, \dots, T_n) \longrightarrow \\
& \quad \text{Time@T}, \mathcal{P}(p_1, \dots, p_n), \text{Dr}(Id, x_b, y_b, E+1)@(T+1) \\
& \text{Time@T}, Pt(p_1, X_1, Y_1)@T_1, \dots, Pt(p_i, X, Y)@T_i, \dots, Pt(p_n, X_n, Y_n)@T_n, \text{Dr}(Id, X, Y, E)@T \\
& \quad \mid \text{doClick}(Id, X, Y, E, T, T_1, \dots, T_i, \dots, T_n) \longrightarrow \\
& \text{Time@T}, Pt(p_1, X_1, Y_1)@T_1, \dots, Pt(p_i, X, Y)@T_i, \dots, Pt(p_n, X_n, Y_n)@T_n, \text{Dr}(Id, X, Y, E-1)@(T+1) \\
& \text{Time@T}, \text{Dr}(Id, X, Y, E)@T \mid \text{hasWind}(X, Y, \text{north}) \longrightarrow \text{Time@T}, \text{Dr}(Id, X, Y+1, E)@(T+1) \\
& \text{Time@T}, \text{Dr}(Id, X, Y+1, E)@T \mid \text{hasWind}(X, Y, \text{south}) \longrightarrow \text{Time@T}, \text{Dr}(Id, X, Y, E)@(T+1) \\
& \text{Time@T}, \text{Dr}(Id, X+1, Y, E)@T \mid \text{hasWind}(X, Y, \text{west}) \longrightarrow \text{Time@T}, \text{Dr}(Id, X, Y, E)@(T+1) \\
& \text{Time@T}, \text{Dr}(Id, X, Y, E)@T \mid \text{hasWind}(X, Y, \text{east}) \longrightarrow \text{Time@T}, \text{Dr}(Id, X+1, Y, E)@(T+1)
\end{aligned}$$

Fig. 1: Macro rules specifying the scenario where drones take pictures of points of interest. Here $\mathcal{P}(p_1, \dots, p_n)$ denotes $P(p_1, X_1, Y_1)@T_1, \dots, P(p_n, X_n, Y_n)@T_n$. Moreover, we assume that the Drone stay in a grid of size $x_{max} \times y_{max}$ and have at most e_{max} energy units.

Proposition 2. *Let \mathcal{A} be a PTS, \mathcal{S}_0 an initial configuration and m the number of facts in \mathcal{S}_0 . For all traces \mathcal{P} of \mathcal{A} starting from \mathcal{S}_0 , let*

$$\mathcal{S}_i \longrightarrow_{\text{Tick}} \mathcal{S}_{i+1} \longrightarrow \dots \longrightarrow \mathcal{S}_j \longrightarrow_{\text{Tick}} \mathcal{S}_{j+1}$$

be any sub-sequence of \mathcal{P} with exactly two instances of the Tick rule, one at the beginning and the other at the end. Then $j - i < m$.

Proposition 3. *Let \mathcal{A} be a PTS. In all infinite traces of \mathcal{A} the global time tends to infinity.*

For readability, we will assume from this point onwards that for all rules r , the set of its constraints implicitly contains the set \mathcal{C}_r as shown in Definition 3, not writing \mathcal{C}_r explicitly in our specifications.

Finally, notice that PTS has many syntactical conditions, *e.g.*, balanced condition (Definition 2), time constraints (Eq. 4), instantaneous rules (Eq. 3). Each one of these conditions have been carefully developed as without any of them important verification problems, such as the reachability problem, becomes undecidable as we show in our previous [13]. Thus these conditions are needed also for infinite traces. The challenge here of allowing infinite traces is to make sure time advances. The definition of PTS is a simple and elegant way to enforce this. Moreover, as we show in Section 3, it is still possible to specify many interesting examples including our motivating example and still prove the decidability of our verification problems involving infinite traces (Section 5).

3 Programming Drone Behavior using PTS

Figure 1 depicts the macro rules of our motivating scenario where drones are moving on a fixed grid of size $x_{max} \times y_{max}$, have at most e_{max} energy units and take pictures of some points of interest. We assume that there are n such points p_1, \dots, p_n , where n

is fixed, a base station is at position (x_b, y_b) , and that the drones should take pictures so that all pictures are recent, that is, the last time a photo of it was taken should not be more than M time units before the current time of any moment.

Clearly if drones choose non-deterministically to move some direction without a particular strategy, they will fail to achieve the assigned goal. A strategy is specified by using time constraints. For this example, the strategy would depend on the difference $T - T_i$, for $1 \leq i \leq n$, specifying the time since the last picture of the point p_i that is the set of time constraints:

$$\mathcal{T}(d_1, \dots, d_n) = \{T - T_1 = d_1, \dots, T - T_n = d_n\}$$

where for all $1 \leq i \leq n$ we instantiate d_i by values in $\{0, \dots, M\}$.

For example, the macro rule with $doMove(Id, X, Y, E + 1, T, T_1, \dots, T_n, north)$ in Figure 1 is replaced by the set of rules:

$$\begin{aligned} &Time@T, \mathcal{P}(p_1, \dots, p_n), Dr(d1, 0, 0, 1)@T \mid \mathcal{T}(0, \dots, 0), DoMov(d1, 0, 0, 1, 0, \dots, 0, north) \longrightarrow \\ &\quad Time@T, \mathcal{P}(p_1, \dots, p_n), Dr(Id, 0, 1, 0)@(T + 1) \\ &Time@T, \mathcal{P}(p_1, \dots, p_n), Dr(d1, 0, 0, 1)@T \mid \mathcal{T}(0, \dots, 1), DoMov(d1, 0, 0, 1, 0, \dots, 1, north) \longrightarrow \\ &\quad Time@T, \mathcal{P}(p_1, \dots, p_n), Dr(Id, 0, 1, 0)@(T + 1) \\ &\dots \\ &Time@T, \mathcal{P}(p_1, \dots, p_n), Dr(d2, x_{max}, y_{max} - 1, e_{max})@T \\ &\quad \mid \mathcal{T}(M, \dots, M), DoMov(d2, x_{max}, y_{max} - 1, e_{max}, M, \dots, M, north) \longrightarrow \\ &\quad Time@T, \mathcal{P}(p_1, \dots, p_n), Dr(Id, x_{max}, y_{max}, e_{max} - 1)@(T + 1) \end{aligned}$$

where $doMove$ is function that returns a boolean value depending on the desired behavior of the drone.

Finally, there are macro rules for moving the drone, taking a picture, charging, and macro specifying winds. While most of the rules have the expected result, we explain the click and wind rules. The click rule is applicable if the drone is at the same position, (X, Y) , as a point of interest p_i . If applied, the timestamp of the fact $P(p_i, X, Y)$ is updated to the current time T . The wind rule is similar to the move rules moving the drone to some direction, but does not cause the drone to consume its energy.

In our implementation, we used a more sophisticated approach described in [24] using soft-constraints to specify a drone's strategy. It can be translated as a PTS by incorporating the strategy used as described above.

Other Examples Finally, there are a number of other examples which we have been investigating and that can be progressive. In [23], we model a simplified version of a package delivery systems inspired by Amazon's Prime Air service. In [24], we model a patrolling bot which moves from one point to another. All these examples seem to be progressive.

Other examples besides those involving drones also seem to be progressive. For example, in our previous work, we specify a monitor for clinical trials [13] using our timed MSR framework with discrete time. This specification seems to be also progressive.

4 Quantitative Temporal Properties

In order to define quantitative temporal properties, we review the notion of critical configurations and compliant traces from our previous work [15]. *Critical configuration*

specification is a set of pairs $\mathcal{CS} = \{\langle \mathcal{S}_1, \mathcal{C}_1 \rangle, \dots, \langle \mathcal{S}_n, \mathcal{C}_n \rangle\}$. Each pair $\langle \mathcal{S}_j, \mathcal{C}_j \rangle$ is of the form:

$$\langle \{F_1 @ T_1, \dots, F_p @ T_p\}, \mathcal{C}_j \rangle$$

where T_1, \dots, T_p are time variables, F_1, \dots, F_p are facts (possibly containing variables) and \mathcal{C}_j is a set of time constraints involving only the variables T_1, \dots, T_p . Given a critical configuration specification, \mathcal{CS} , we classify a configuration \mathcal{S} as *critical* if for some $1 \leq i \leq n$, there is a grounding substitution, σ , mapping time variables in \mathcal{S}_i to natural numbers and non time variables to terms such that:

- $\mathcal{S}_i \sigma \subseteq \mathcal{S}$;
- all constraints in \mathcal{C}_i are valid.

where substitution application ($\mathcal{S}\sigma$) is defined as usual [9].

Example 1. We can specify usual safety conditions which do not involve time. For example, a drone should never run out of energy. This can be specified by using the following set of critical configuration specification:

$$\{\langle \{Dr(Id, X, Y, 0) @ T\}, \emptyset \rangle \mid Id \in \{d1, d2\}, X \in \{0, \dots, x_{max}\}, Y \in \{0, \dots, y_{max}\}\}$$

Example 2. The following critical configuration specification specifies a quantitative property involving time:

$$\{\langle \{P(p_1, x_1, y_1) @ T_1, Time @ T\}, T > T_1 + M \rangle, \dots, \langle \{P(p_n, x_n, y_n) @ T_n, Time @ T\}, T > T_n + M \rangle \}$$

Together with the specification in Figure 1, this critical configuration specification specifies that the last pictures of all points of interest (p_1, \dots, p_n located at $(x_1, y_1), \dots, (x_n, y_n)$) should have timestamps no more than M time units old.

Example 3. Let the facts $St(Id) @ T_1$ and $St(empty) @ T_1$ denote, respectively, that at time T_1 the drone Id entered the base station to recharge and that the station is empty. Moreover, assume that only one drone may be in the station to recharge, which would be specified by adding the following rules specifying the drone landing and take off, where st is a constant symbol denoting that a drone landed on the base station:

$$\begin{aligned} Time @ T, Dr(Id, x_b, y_b) @ T, St(empty) @ T_1 &\longrightarrow Time @ T, Dr(Id, st, st) @ (T + 1), St(Id) @ T \\ Time @ T, Dr(Id, st, st) @ T, St(Id) @ T_1 &\longrightarrow Time @ T, Dr(Id, x_b, y_b) @ (T + 1), St(empty) @ T \end{aligned}$$

Then, the critical configuration specification $\{\langle \{St(Id) @ T_1, Time @ T\}, T > T_1 + M_1 \rangle \mid Id \in \{d1, d2\}\}$ specifies that one drone should not remain too long (more than M_1 time units) in a base station not allowing other drones to charge.

Definition 4. A trace of a timed MSR is compliant for a given critical configuration specification if it does not contain any critical configuration.

We will be interested in survivability which requires checking whether, given an initial configuration, all possible infinite traces of a system are compliant. In order to define a sensible notion of survivability, however, we need to assume some conditions on when the Tick rule is applicable. With no conditions on the application of the Tick rule many timed systems of interest, such as our main example with drones, do not satisfy survivability as the following trace containing only instances of the Tick rule could always be constructed:

$$\mathcal{S}_1 \longrightarrow_{Tick} \mathcal{S}_2 \longrightarrow_{Tick} \mathcal{S}_3 \longrightarrow_{Tick} \mathcal{S}_4 \longrightarrow_{Tick} \dots$$

Imposing a *time sampling* is a way to avoid such traces where the time simply ticks. They are used, for example, in the semantics of verification tools such as Real-Time Maude [20]. In particular, a time sampling dictates when the Tick rule must be applied and when it cannot be applied. This treatment of time is used both for dense and discrete times in searching and model checking timed systems.

Definition 5. A (possibly infinite) trace \mathcal{P} of a timed MSR \mathcal{A} uses a lazy time sampling if for any occurrence of the Tick rule $\mathcal{S}_i \longrightarrow_{Tick} \mathcal{S}_{i+1}$ in \mathcal{P} , no instance of any instantaneous rule in \mathcal{A} can be applied to the configuration \mathcal{S}_i .

In lazy time sampling instantaneous rules are given a higher priority than the Tick rule. Under this time sampling, a drone should carry out one of the rules in Figure 1 at each time while time can only advance when all drones have carried out their actions for that moment. This does not mean, however, that the drones will satisfy their goal of always having recent pictures of the points of interest as this would depend on the behavior of the system, *i.e.*, the actions carried out by the drones. Intuitively, the lazy time sampling does not allow the passing of time if there are scheduled drone actions at the current time. Its semantics reflects that all undertaken actions do happen.

In the remainder of this paper, we fix the time sampling to lazy time sampling. We leave for future work investigating whether our complexity results hold for other time samplings.

4.1 Verification Problems

The first property we introduce is realizability. Realizability is useful for increasing one's confidence in a specified system, as clearly a system that is not realizable can not accomplish the given tasks (specified by a critical specification) and therefore, the designer would need to reformulate it. However, if a system is shown realizable, the trace, \mathcal{P} , used to prove it could also provide insights on the sequence of actions that lead to accomplishing the specified tasks. This may be used to refine the specification reducing possible non-determinism.

Definition 6. A timed MSR \mathcal{A} is realizable (*resp.*, n -time-bounded realizable) with respect to the lazy time sampling, a critical configuration specification \mathcal{CS} and an initial configuration \mathcal{S}_0 if there exists a trace, \mathcal{P} , that starts with \mathcal{S}_0 and uses the lazy time sampling such that:

1. \mathcal{P} is compliant with respect to \mathcal{CS} ;
2. Global time tends to infinity (*resp.*, global time advances by exactly n time units) in \mathcal{P} .

The second condition that global time tends to infinity, which implies that only a finite number of actions are performed in a given time. Another way of interpreting this condition following [1] is of a liveness condition, that is, the system should not get stuck. The first condition, on the other hand, is a safety condition as it states that no bad state should be reached. Thus the feasibility problem (and also the survivability problem introduced next) is a combination of a liveness and safety conditions. Moreover, since \mathcal{CS} involve time constraints, it is a quantitative liveness and safety property.

The n -time-bounded realizability problem is motivated by bounded model checking. We look for a finite compliant trace that spreads over a n units of time, where n is fixed.

As already noted, realizability could be useful in reducing non-determinism in the specification. In many cases, however, it is not desirable and even not possible to eliminate the non-determinism of the system. For example, in open distributed systems, the environment can play an important role. Winds, for example, may affect drones' performances such as the speed and energy required to move from one point to another. We would like to know whether for all possible decisions taken by agents and under the interference of the environment, the given timed MSR accomplishes the specified tasks. *If so, we say that a system satisfies survivability.*

Definition 7. A timed MSR \mathcal{A} satisfies survivability (resp., n -time-bounded survivability) with respect to the lazy time sampling, a critical configuration specification \mathcal{CS} and an initial configuration \mathcal{S}_0 if it is realizable (resp., n -time-bounded realizable) and if all infinite traces (resp. all traces with exactly n instances of the Tick rule), \mathcal{P} , that start with \mathcal{S}_0 and use the lazy time sampling are such that:

1. \mathcal{P} is compliant with respect to \mathcal{CS} ;
2. The global time tends to infinity (resp., no condition).

5 Complexity Results

Our complexity results, for a given PTS \mathcal{A} , an initial configuration \mathcal{S}_0 and a critical configuration specification \mathcal{CS} , will mention the value D_{max} which is an upper-bound on the natural numbers appearing in \mathcal{S}_0 , \mathcal{A} and \mathcal{CS} . D_{max} can be inferred syntactically by simply inspecting the timestamps of \mathcal{S}_0 , the D values in timestamps of rules (which are of the form $T + D$) and constraints in \mathcal{A} and \mathcal{CS} (which are of the form $T_1 > T_2 + D$ and $T_1 = T_2 + D$). For example, the $D_{max} = 1$ for the specification in Figure 1.

The size of a timestamped fact $F@T$, written $|F@T|$ is the total number of alphabet symbols appearing in F . For instance, $|P(s(z), f(a, X), a)@12| = 7$. For our complexity results, we assume a bound, k , on the size of facts. For example, in our specification in Figure 1, we can take the bound $k = |x_{max}| + |y_{max}| + |e_{max}| + 5$. Without this bound (or other restrictions), any interesting decision problem is undecidable by encoding the Post correspondence problem [8].

Notice that we do not always impose an upper bound on the values of timestamps.

Assume throughout this section the following: (1) Σ – A finite alphabet with J predicate symbols and E constant and function symbols; \mathcal{A} – A PTS constructed over Σ ; m – The number of facts in the initial configuration \mathcal{S}_0 ; \mathcal{CS} – A critical configuration specification constructed over Σ ; k – An upper-bound on the size of facts; D_{max} – An upper-bound on the numeric values of \mathcal{S}_0 , \mathcal{A} and \mathcal{CS} .

5.1 PSPACE-Completeness

In order to prove the PSPACE-completeness of realizability and survivability problems, we review the machinery introduced in our previous work [13] called δ -configuration.

For a given D_{max} the *truncated time difference* of two timed facts $P@t_1$ and $Q@t_2$ with $t_1 \leq t_2$, denoted by $\delta_{P,Q}$, is defined as follows:

$$\delta_{P,Q} = \begin{cases} t_2 - t_1, & \text{provided } t_2 - t_1 \leq D_{max} \\ \infty, & \text{otherwise} \end{cases}$$

Let $\mathcal{S} = Q_1 @ t_1, Q_2 @ t_2, \dots, Q_n @ t_n$, be a configuration of a timed MSR \mathcal{A} written in canonical way where the sequence of timestamps t_1, \dots, t_n is non-decreasing. The δ -configuration of \mathcal{S} for a given D_{max} is

$$\delta_{\mathcal{S}, D_{max}} = [Q_1, \delta_{Q_1, Q_2}, Q_2, \dots, Q_{n-1}, \delta_{Q_{n-1}, Q_n}, Q_n].$$

In our previous work [15,13], we showed that a δ -configuration is an equivalence class on configurations. Namely, for a given D_{max} , we declare \mathcal{S}_1 and \mathcal{S}_2 equivalent, written $\mathcal{S}_1 \equiv_{D_{max}} \mathcal{S}_2$, if and only if their δ -configurations are exactly the same. Moreover, we showed that there is a bisimulation between (compliant) traces over configurations and (compliant) traces over their δ -configurations in the following sense: if $\mathcal{S}_1 \longrightarrow \mathcal{S}_2$ and $\mathcal{S}_1 \equiv_{D_{max}} \mathcal{S}'_1$, then there is a trace $\mathcal{S}'_1 \longrightarrow \mathcal{S}'_2$ such that $\mathcal{S}_2 \equiv_{D_{max}} \mathcal{S}'_2$. This result appears in [15, Corollary 7] and more details can be found in Appendix A.

Therefore, in the case of balanced timed MSRs, we can work on traces constructed using δ -configurations. Moreover, the following lemma establishes a bound on the number of different δ -configurations. The proof can be found in Appendix B.

Lemma 1. *Assume $\Sigma, \mathcal{A}, \mathcal{S}_0, m, \mathcal{CS}, k, D_{max}$ as described above. The number of different δ -configurations, denoted by $L_\Sigma(m, k, D_{max})$ is such that*

$$L_\Sigma(m, k, D_{max}) \leq (D_{max} + 2)^{(m-1)} J^m(E + 2mk)^{mk}.$$

Infinite Traces Our previous work only dealt with *finite traces*. The challenge here is to deal with infinite traces and in particular the feasibility and survivability problems. These problems are new and as far as we know have not been investigated in the literature (see Section 7 for more details).

PSPACE-hardness of both the realizability and survivability can be shown by adequately adapting our previous work [16] (shown in the Appendix C). We therefore show PSPACE-membership of these problems.

Recall that a system is realizable if there is a compliant infinite trace \mathcal{P} in which the global time tends to infinity. Since \mathcal{A} is progressive, we get the condition on time from Proposition 3. We, therefore, need to construct a compliant infinite trace. The following lemma establishes a criteria:

Lemma 2. *Assume $\Sigma, \mathcal{A}, \mathcal{S}_0, m, \mathcal{CS}, k, D_{max}$ as described above. If there is a compliant trace (constructed using δ -configurations) starting with (the δ -representation of) \mathcal{S}_0 with length $L_\Sigma(m, k, D_{max})$, then there is an infinite compliant trace starting with (the δ -representation of) \mathcal{S}_0 .*

Assume that for any given timed MSR \mathcal{A} , an initial configuration \mathcal{S}_0 and a critical configuration specification \mathcal{CS} we have two functions \mathcal{N} and \mathcal{X} which check, respectively, whether a rule in \mathcal{A} is applicable to a given δ -configuration and whether a δ -configuration is critical with respect to \mathcal{CS} . Moreover, let \mathcal{T} be a function implementing the lazy time sampling. It takes a timed MSR and a δ -configuration of that system, and returns 1 when the tick must be applied and 0 when it must not be applied. We assume that \mathcal{N} , \mathcal{X} and \mathcal{T} run in Turing time bounded by a polynomial in $m, k, \log_2(D_{max})$. Notice that for our examples this is the case. Because of Lemma 2, we can show that the realizability

problem is in PSPACE by searching for compliant traces of length $L_\Sigma(m, k, D_{max})$ (stored in binary). To do so, we rely on the fact that PSPACE and NPSPACE are the same complexity class [22].

Theorem 1. *Assume Σ a finite alphabet, \mathcal{A} a PTS, an initial configuration S_0 , m the number of facts in S_0 , CS a critical configuration specification, k an upper-bound on the size of facts, D_{max} an upper-bound on the numeric values in S_0 , \mathcal{A} and CS , and the functions \mathcal{N} , \mathcal{X} and \mathcal{T} as described above. There is an algorithm that, given an initial configuration S_0 , decides whether \mathcal{A} is realizable with respect to the lazy time sampling, CS and S_0 and the algorithm runs in space bounded by a polynomial in m , k and $\log_2(D_{max})$.*

The polynomial is in fact $\log_2(L_\Sigma(m, k, D_{max}))$ and the proof is in Appendix D.

We now consider the survivability problem. Recall that in order to prove that \mathcal{A} satisfies survivability with respect to the lazy time sampling, CS and S_0 , we must show that \mathcal{A} is realizable and that for all infinite traces \mathcal{P} starting with S_0 (Definition 7):

1. \mathcal{P} is compliant with respect to CS ;
2. The global time in \mathcal{P} tends to infinity.

Checking that a system is realizable is PSPACE-complete as we have just shown. Moreover, the second property (time tends to infinity) follows from Proposition 3 for progressive timed MSR. It remains to show that all infinite traces using the lazy time sampling are compliant, which reduces to checking that *no critical configuration is reachable* from the initial configuration S_0 by a trace using the lazy time sampling. This property can be decided in PSPACE by relying on the fact that PSPACE, NPSPACE and co-PSPACE are all the same complexity class [22]. Therefore, survivability is also in PSPACE as states the following theorem. Its proof can be found in Appendix E.

Theorem 2. *Assume Σ , \mathcal{A} , S_0 , m , CS , k , D_{max} and the functions \mathcal{N} , \mathcal{X} and \mathcal{T} as described in Theorem 1. There is an algorithm that decides whether \mathcal{A} satisfies the survivability problem with respect to the lazy time sampling, CS and S_0 which runs in space bounded by a polynomial in m , k and $\log_2(D_{max})$.*

Corollary 1. *Both the realizability and the survivability problem for PTS are PSPACE-complete when assuming a bound on the size of facts.*

5.2 Complexity Results for n -Time-Bounded Systems

We now consider the n -time-bounded versions of the Realizability and Survivability problems (Definitions 6 and 7).

The following lemma establishes an upper-bound on the length of traces with exactly n instances of tick rules for PTS. It follows immediately from Proposition 2.

Lemma 3. *Let n be fixed and assume Σ , \mathcal{A} , S_0 , m , CS , k , D_{max} as described in Theorem 1. For all traces \mathcal{P} of \mathcal{A} with exactly n instances of the Tick rule, the length of \mathcal{P} is bounded by $(n + 2) * m + n$.*

We can check in polynomial time whether a trace is compliant and has exactly n Ticks. Therefore, the n -time-bounded realizability problem is in NP as stated by the following theorem. Its proof is in the Appendix F.

Exp 1: ($N = 1, P = 4, x_{max} = y_{max} = 10$)		Exp 3: ($N = 2, P = 9, x_{max} = y_{max} = 20$)	
$M = 50, e_{max} = 40$	F, $st = 139, t = 0.3$	$M = 100, e_{max} = 500$	F, $st = 501, t = 6.2$
$M = 70, e_{max} = 40$	F, $st = 203, t = 0.4$	$M = 150, e_{max} = 500$	F, $st = 1785, t = 29.9$
$M = 90, e_{max} = 40$	S, $st = 955, t = 2.3$	$M = 180, e_{max} = 500$	S, $st = 2901, t = 49.9$
		$M = 180, e_{max} = 150$	F, $st = 1633, t = 25.6$
Exp 2: ($N = 2, P = 4, x_{max} = y_{max} = 10$)		Exp 4: ($N = 3, P = 9, x_{max} = y_{max} = 20$)	
$M = 30, e_{max} = 40$	F, $st = 757, t = 3.2$	$M = 100, e_{max} = 150$	F, $st = 3217, t = 71.3$
$M = 40, e_{max} = 40$	F, $st = 389, t = 1.4$	$M = 120, e_{max} = 150$	F, $st = 2193, t = 52.9$
$M = 50, e_{max} = 40$	S, $st = 821, t = 3.2$	$M = 180, e_{max} = 150$	S, $st = 2193, t = 53.0$
		$M = 180, e_{max} = 100$	F, $st = 2181, t = 50.4$

Table 1: N is the number of drones, P the number of points of interest, $x_{max} \times y_{max}$ the size of the grid, M the time limit for photos, and e_{max} the maximum energy capacity of each drone. We measured st and t , which are, respectively, the number of states and time in seconds until finding a counter example if F (fail), and until searching all traces with exactly $4 \times M$ ticks if S (success).

Theorem 3. *Let n be fixed and assume $\Sigma, \mathcal{A}, \mathcal{S}_0, m, \mathcal{CS}, k, D_{max}$ and the functions $\mathcal{N}, \mathcal{X}, \mathcal{T}$ as described in Theorem 1. The problem of determining whether \mathcal{A} is n -time-bounded realizable with respect to the lazy time sampling, \mathcal{CS} and \mathcal{S}_0 is in NP with \mathcal{S}_0 as the input.*

For NP-hardness, we encode the NP-hard problem 3-SAT as an n -time-bounded realizability problem as done in our previous work [14]. The encoding can be found in the Appendix G.

Recall that for n -time-bounded survivability property, we need to show that:

1. \mathcal{A} is n -time-bounded realizable with respect to \mathcal{CS} ;
2. All traces using the lazy time sampling with exactly n ticks are compliant with respect to \mathcal{CS} .

As we have shown, the first sub-problem is NP-complete. The second sub-problem is reduced to checking that no critical configuration is reachable from \mathcal{S}_0 by a trace using the lazy time sampling with less or equal to n ticks. We do so by checking whether a critical configuration is reachable. This is similar to realizability which we proved to be in NP. If a critical configuration is reachable then \mathcal{A} does not satisfy the second sub-problem, otherwise it does satisfy. Therefore, deciding the second sub-problem is in co-NP. Thus the n -timed survivability problem is in a class containing both NP and co-NP, e.g., Δ_2^P of the polynomial hierarchy (P^{NP}) [21].

Theorem 4. *Let n be fixed and assume $\Sigma, \mathcal{A}, \mathcal{S}_0, m, \mathcal{CS}, k, D_{max}$ and the functions $\mathcal{N}, \mathcal{X}, \mathcal{T}$ as described in Theorem 1. The problem of determining whether \mathcal{A} satisfies n -time-bounded survivability with respect to the lazy time sampling, \mathcal{CS} and \mathcal{S}_0 is in the class Δ_2^P of the polynomial hierarchy (P^{NP}) with input \mathcal{S}_0 .*

6 Bounded Simulations

For our bounded simulations, we implemented a more elaborated version of our running scenario in Maude using the machinery described in [24]. Our preliminary results are very promising. We are able to model-check fairly large systems for the bounded survivability.

We consider N drones which should have recent pictures, *i.e.*, at most M time units old, of P points distributed in a grid $x_{max} \times y_{max}$, where the base station is at position $(\lceil x_{max}/2 \rceil, \lceil y_{max}/2 \rceil)$, and drones have maximum energy of e_{max} . Drones use soft-constraints, which take into account the drone's position, energy, and pictures, to rank their actions and they perform any one the best ranked actions. Drones are also able to share information with the base station.

Our simulation results are depicted in Table 1. We model-checked the n -timed survivability of the system where $n = 4 \times M$. We varied M and the maximum energy capacity of drones e_{max} . Our implementation [24] finds counter examples quickly (less than a minute) even when considering a larger grid (20×20) and three drones.⁹

We can observe that our implementations can help specifiers to decide how many drones to use and with which energy capacities. For example, in Exp 3, drones required a great deal of energy, namely 500 energy units. Adding an additional drone, Exp 4, reduced the energy needed to 150 energy units. Finally, the number of states may increase when decreasing M because with lower values of M , drones may need to come back more often to the base station causing them to share information and increasing the number of states.

7 Related and Future Work

This paper introduced a novel sub-class of timed MSR systems called progressive which is defined by imposing syntactic restrictions on MSR rules. We illustrated with examples of Time Sensitive Distributed Systems that this is a relevant class of systems. We also introduced two verification problems which may depend on explicit time constraints, namely realizability and survivability, defined over infinite traces. We showed that both problems are PSPACE-complete for progressive timed systems, and when we additionally impose a bound on time, realizability becomes NP-complete and survivability is in Δ_2^P of the polynomial hierarchy. Finally, we demonstrated by experiments that it is feasible to analyse fairly large progressive systems using the rewriting logic tool Maude.

Others have proposed languages for specifying properties which allow explicit time constraint. We review some of the timed automata, temporal logic and rewriting literature.

Our progressive condition is related to the *finite-variability assumption* used in the temporal logic and timed automata literature [10,17,18,2,3]: in any bounded interval of time, there can be only finitely many observable events or state changes. Similarly, progressive systems have the property that only a finite number of instantaneous rules can be applied in any bounded interval of time (Proposition 2). Such a property seems necessary for the decidability of many temporal verification problems.

⁹ Although these scenarios seem small, the state space grow very fast: the state space of our largest scenario has an upper bound of $(400 \times 399 \times 398) \times (150 \times 150 \times 150) \times (180 \times 4) \times (180)^9 \geq 3.06 \times 10^{37}$ states.

As we discussed in much more detail in the Related Work section of our previous work [13], there are some important differences between our timed MSR and timed automata [2,3] on both the expressive power and decidability proofs. For example, a description of a timed MSR system uses first order formulas with variables, whereas timed automata are able to refer only to transition on ground states. That is, timed MSR is essentially a first-order language, while timed automata are propositional. If we replace a first order description of timed MSR by all its instantiations, that would lead to an exponential explosion. Furthermore, in contrast with the timed automata paradigm, in timed MSR we can manipulate in a natural way the facts both in the past, in the future, and in the present. Finally, our model uses discrete times, while timed automata uses dense times. It seems, however, possible to extend our results to dense times given our previous work [12]. We leave this investigation to future work.

The temporal logic literature has proposed many languages for the specification and verification of timed systems. While many temporal logics include quantitative temporal operators, *e.g.* [18,17], this literature does not discuss notions similar to realizability and survivability notions introduced here. In addition to that, our specifications are executable. Indeed, as we have done here, our specifications can be executed in Maude.

The work [1,5] classifies traces and sets of traces as safety, liveness or properties that can be reduced to subproblems of safety and liveness. Following this terminology, properties relating to both of our problems of realizability and survivability (that involve infinite traces) contain elements of safety as well as elements of liveness. Properties relating to the n -time-bounded versions of realizability and survivability could be classified as safety properties. We do not see how to express this in the terms of [1,5]. We intend to revisit this in future work.

Real-Time Maude is a tool for simulating and analyzing real-time systems. Rewrite rules are partitioned into instantaneous rules and rules that advance time, where instantaneous rules are given priority. Time advance rules may place a bound on the amount of time to advance, but do not determine a specific amount, thus allowing continual observation of the system. Time sampling strategies are used to implement search and model-checking analyses. Ölveczky and Messeguer [19] investigate conditions under which the maximal time sampling strategy used in Real-Time Maude is complete. One of the conditions required is tick-stabilizing which is similar to progressive and the finite variability assumption in that one assumes a bound on the number of actions applicable in a finite time.

Cardenas *et al.* [4] discuss possible verification problems of cyber-physical systems in the presence of malicious intruders. They discuss surviving attacks, such as denial of service attacks on the control mechanisms of devices. We believe that our progressive timed systems can be used to define sensible intruder models and formalize the corresponding survivability notions. This may lead to the automated analysis of such systems similar to the successful use of the Dolev-Yao intruder model [7] for protocol security verification. Given the results of this paper, for the decidability of any security problem would very likely involve a progressive timed intruder model.

Finally, we believe it is possible to extend this work to dense times given our previous work [12]. There we assume a Tick rule of the form $Time@T \longrightarrow Time@(T + \epsilon)$.

However, we do not consider critical configuration specifications. We are currently investigating how to incorporate the results in this paper with the results of [12].

References

1. Bowen Alpern and Fred B. Schneider. Recognizing safety and liveness. *Distributed Computing*, 2(3):117–126, 1987.
2. Rajeev Alur and Thomas A. Henzinger. Logics and models of real time: A survey. In *REX Workshop*, pages 74–106, 1991.
3. Rajeev Alur and P. Madhusudan. Decision problems for timed automata: A survey. In *SFM*, pages 1–24, 2004.
4. Alvaro A. Cárdenas, Saurabh Amin, and Shankar Sastry. Secure control: Towards survivable cyber-physical systems. In *ICDCS*, pages 495–500, 2008.
5. Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
6. Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn Talcott. *All About Maude: A High-Performance Logical Framework*, 2007.
7. D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
8. Nancy A. Durgin, Patrick Lincoln, John C. Mitchell, and Andre Scedrov. Multiset rewriting and the complexity of bounded security protocols. *Journal of Computer Security*, 12(2):247–311, 2004.
9. Herbert B. Enderton. *A mathematical introduction to logic*. Academic Press, 1972.
10. Marco Faella, Axel Legay, and Mariëlle Stoelinga. Model checking quantitative linear time logic. *Electr. Notes Theor. Comput. Sci.*, 220(3):61–77, 2008.
11. Max Kanovich, Tajana Ban Kirigin, Vivek Nigam, and Andre Scedrov. Bounded memory Dolev-Yao adversaries in collaborative systems. *Inf. Comput.*, 2014.
12. Max Kanovich, Tajana Ban Kirigin, Vivek Nigam, Andre Scedrov, and Carolyn Talcott. Discrete vs. dense times in the analysis of cyber-physical security protocols. In *POST*, pages 259–279, 2015.
13. Max Kanovich, Tajana Ban Kirigin, Vivek Nigam, Andre Scedrov, and Carolyn Talcott. A rewriting framework and logic for activities subject to regulations. *Mathematical Structures in Computer Science*, 2015. Published online.
14. Max I. Kanovich, Tajana Ban Kirigin, Vivek Nigam, and Andre Scedrov. Bounded memory protocols and progressing collaborative systems. In *ESORICS*, 2013.
15. Max I. Kanovich, Tajana Ban Kirigin, Vivek Nigam, Andre Scedrov, Carolyn L. Talcott, and Ranko Perovic. A rewriting framework for activities subject to regulations. In *RTA*, pages 305–322, 2012.
16. Max I. Kanovich, Paul Rowe, and Andre Scedrov. Collaborative planning with confidentiality. *J. Autom. Reasoning*, 46(3-4):389–421, 2011.
17. François Laroussinie, Philippe Schnoebelen, and Mathieu Turuani. On the expressivity and complexity of quantitative branching-time temporal logics. *Theor. Comput. Sci.*, 2003.
18. Carsten Lutz, Dirk Walther, and Frank Wolter. Quantitative temporal logics: PSPACE and below. In *TIME*, pages 138–146, 2005.
19. Peter Csaba Ölveczky and José Meseguer. Abstraction and completeness for real-time maude. *Electr. Notes Theor. Comput. Sci.*, 176(4):5–27, 2007.
20. Peter Csaba Ölveczky and José Meseguer. The real-time maude tool. In *TACAS 2008*, pages 332–336, 2008.

21. Christos H. Papadimitriou. *Computational complexity*. Academic Internet Publ., 2007.
22. W. J. Savitch. Relationship between nondeterministic and deterministic tape classes. *Journal of Computer and System Sciences*, 4:177–192, 1970.
23. Carolyn L. Talcott, Farhad Arbab, and Maneesh Yadav. Soft agents: Exploring soft constraints to model robust adaptive distributed cyber-physical agent systems. In *Software, Services, and Systems - Essays Dedicated to Martin Wirsing*, pages 273–290, 2015.
24. Carolyn Talcott, Vivek Nigam, Farhad Arbab, and Tobias Kappé. Formal specification and analysis of robust adaptive distributed cyber-physical systems. In *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems* 2016.

Acknowledgments: Kanovich’s research was partially supported by EPSRC. Scedrov’s research was partially supported by ONR and by AFOSR MURI. Kanovich’s and Scedrov’s work on this paper was partially carried out within the framework of the Basic Research Program at the National Research University Higher School of Economics (HSE) and partially supported within the framework of a subsidy by the Russian Academic Excellence Project ‘5-100’. Talcott was partially supported by NSF grant CNS-1318848 and ONR grant N00014-15-1-2202. Nigam and Talcott were partially supported by Capes Science without Borders grant 88881.030357/2013-01. Nigam was partially supported by Capes and CNPq.

A Bisimulation result

Theorem 5. *For any timed MSR \mathcal{A} the equivalence relation between configurations is well-defined with respect to the actions of the system (including time advances), lazy time scheduling and critical configurations. Any compliant trace starting from the given initial configuration can be conceived as a compliant trace over δ -representations.*

Proof. The equivalence among configurations is well defined with respect to application of actions, *i.e.* action application on δ -representations is independent of the choice of configuration from the same class. More precisely, as shown in the diagram below, assume \mathcal{S}_1 and \mathcal{S}_2 are equivalent configurations, and assume that \mathcal{S}_1 is transformed to \mathcal{S}'_1 by means of an action α . Notice that equivalent configurations satisfy the same set of constraints. Hence, the action α is applicable to \mathcal{S}_2 and will transform \mathcal{S}_2 into some \mathcal{S}'_2 :

$$\begin{array}{c} \mathcal{S}_1 \rightarrow_{\alpha} \mathcal{S}'_1 \\ \wr \\ \mathcal{S}_2 \rightarrow_{\alpha} \mathcal{S}'_2 \end{array}$$

It remains to show that \mathcal{S}'_1 is equivalent to \mathcal{S}'_2 . We consider the two types of actions, namely, time advances and instantaneous actions. Notice that, using the lazy time sampling, tick rule is applied to \mathcal{S}_1 if and only if no instantaneous action can be applied to the given configuration \mathcal{S}_1 . Since \mathcal{S}_1 and \mathcal{S}_2 satisfy the same set of constraints, it follows that the tick rule is applied to \mathcal{S}_2 if and only if the tick rule is applied \mathcal{S}_1 .

Let the time advance transform \mathcal{S}_1 into \mathcal{S}'_1 , and \mathcal{S}_2 to \mathcal{S}'_2 . Since only the timestamp T denoting the global time in $Time@T$ is increased by 1, and the rest of the configuration remains unchanged, only truncated time differences involving $Time$ change in the resulting configurations. Because of the equivalence $\mathcal{S}_1 \sim \mathcal{S}_2$, for a fact $P@T_1^P$ in \mathcal{S}_1 with $T_1^P \leq T$, $Time@T$ and $\delta_{P,Time} = t$, we have $P@T_2^P$ with $T_2^P \leq \hat{T}$, $Time@\hat{T}$ and $\delta_{P,Time} = t$ in \mathcal{S}_2 as well. Therefore, we have $\delta_{P,Time} = [t + 1]$ both in \mathcal{S}'_1 and \mathcal{S}'_2 . On the other hand for any future fact $Q@T^Q$ with $\delta_{Time,Q} = t$ in \mathcal{S}_1 and in \mathcal{S}_2 , we get $\delta_{Time,Q} = t - 1$ in both \mathcal{S}'_1 and \mathcal{S}'_2 . Therefore, \mathcal{S}'_1 and \mathcal{S}'_2 are equivalent.

The reasoning the instantaneous actions is similar. Each created fact in \mathcal{S}'_1 and \mathcal{S}'_2 is of the form $P@(T^1 + d)$ and $P@(T^2 + d)$, where T^1 and T^2 represent global time in \mathcal{S}_1 and \mathcal{S}_2 , respectively. Therefore each created fact has the same difference d to the global time in the corresponding configuration. This implies that the created facts have the same truncated time differences to the remaining facts. Hence \mathcal{S}'_1 and \mathcal{S}'_2 are equivalent. Therefore, action application on δ -representations shown in is well defined.

Since equivalent configurations satisfy the same set of constraints, \mathcal{S}_1 is a critical configuration if and only if \mathcal{S}_2 is a critical configuration. By induction on the length of the trace, it immediately follows that, given a timed MSR, any compliant trace over configurations can be represented by a compliant trace over δ -representations. That is, the abstraction of configurations to δ -representations is complete.

The abstraction is also sound. Namely, from a compliant trace over δ -representations, we can extract a concrete compliant trace over configurations. Although any given δ -representation corresponds to an infinite number of configurations, for a given initial configuration \mathcal{S}_0 , we have $\Delta_0 = \delta_{\mathcal{S}_0}$. The existence of a compliant trace over configurations is then easily proven by induction on the length of the trace over δ -representations.

B Bound on the number of different δ -configurations (Lemma 1)

Proof. Let the given finite alphabet contain J predicate symbols and E constant and function symbols. Let the initial configuration S_0 contain m facts. Let

$$[Q_1, \delta_{Q_1, Q_2}, Q_2, \dots, Q_{m-1}, \delta_{Q_{m-1}, Q_m}, Q_m]$$

be a δ -representation with m facts. There are m slots for predicate names and at most mk slots for constants and function symbols, where k is the bound on the size of facts. Constants can be either constants in the initial alphabet or names for fresh values (nonces). Following [11], we need to consider only $2mk$ names for fresh values (nonces). Finally, only time differences up to D_{max} have to be considered together with the symbol ∞ and there are $m - 1$ slots for time differences in a δ -representation. Therefore the number of different δ -configurations is bounded by $(D_{max} + 2)^{(m-1)} J^m (E + 2mk)^{mk}$.

C Encoding of a Turing Machine that accepts in Space n

We encode a non-deterministic Turing machine M that accepts in space n . We adapt the encoding in [11] to a progressive timed MSR \mathcal{A} that uses the lazy time sampling. For readability, in the rules below, we do not explicitly write the set of constraints \mathcal{C}_r as per Definition 3. This set is implicitly assumed.

First, we introduce the following propositions: $R_{i,\xi}@T_l$ which denotes that “the i -th cell contains symbol ξ since time T_l ”, where $i = 0, 1, \dots, n+1$, ξ is a symbol of the tape alphabet of M , and $S_{j,q}$ denotes that “the j -th cell is scanned by M in state q ”, where $j = 0, 1, \dots, n+1$, q is a state of M .

A Turing machine configuration will be encoded by using the multiset of facts:

$$Time@T, S_{j,q}@T_1, R_{0,\xi_0}@T_2, R_{1,\xi_1}@T_2, R_{2,\xi_2}@T_3, \dots, R_{n,\xi_n}@T_{n+2}, R_{n+1,\xi_{n+1}}@T_{n+3}. \quad (5)$$

Second, each instruction γ in M of the form $q\xi \rightarrow q'\eta D$, denoting “if in state q looking at symbol ξ , replace it by η , move the tape head one cell in direction D along the tape, and go into state q' ”, is specified by the set of $5(n+2)$ actions of the form:

$$\begin{aligned} Time@T, S_{i,q}@T, R_{i,\xi}@T &\rightarrow Time@T, F_{i,\gamma}@T, R_{i,\xi}@T \\ Time@T, F_{i,\gamma}@T, R_{i,\xi}@T &\rightarrow Time@T, F_{i,\gamma}@T, H_{i,\gamma}@T \\ Time@T, F_{i,\gamma}@T, H_{i,\gamma}@T &\rightarrow Time@T, G_{i,\gamma}@T, H_{i,\gamma}@T \\ Time@T, G_{i,\gamma}@T, H_{i,\gamma}@T &\rightarrow Time@T, G_{i,\gamma}@T, R_{i,\eta}@T \\ Time@T, G_{i,\gamma}@T, R_{i,\eta}@T &\rightarrow Time@T, S_{i_D,q'}@T(T+1), R_{i,\eta}@T(T+1), \end{aligned} \quad (6)$$

where $i = 0, 1, \dots, n+1$, $F_{i,\gamma}$, $G_{i,\gamma}$, $H_{i,\gamma}$ are auxiliary atomic propositions, $i_D := i+1$ if D is right, $i_D := i-1$ if D is left, and $i_D := i$, otherwise. It is easy to check that above rules are necessarily applied in succession, i.e. the only transition possible is of the following form:

$$\begin{aligned} Time@T, S_{i,q}@T, R_{i,\xi}@T &\rightarrow Time@T, F_{i,\gamma}@T, R_{i,\xi}@T \rightarrow \\ &\rightarrow Time@T, F_{i,\gamma}@T, H_{i,\gamma}@T \rightarrow Time@T, G_{i,\gamma}@T, H_{i,\gamma}@T \rightarrow \\ &\rightarrow Time@T, G_{i,\gamma}@T, R_{i,\eta}@T \rightarrow Time@T, S_{i_D,q'}@T(T+1), R_{i,\eta}@T(T+1). \end{aligned} \quad (7)$$

At this point, no instantaneous rule is applicable and therefore the Tick rule should be applied. Thus the encoding reflects the lazy time sampling.

The critical configuration specification is any configuration corresponding to a final state of the Turing Machine, that is:

$$\{\langle \{S_{i_D, q_F} @ T\}, \emptyset \rangle \mid q_F \text{ is an accepting or rejecting state} \}.$$

By the above encoding we reduce the problem of a Turing machine termination in n -space to the realizability problem. More precisely, the given Turing machine M does not terminate if and only if there is an infinite compliant trace in the obtained progressive timed MSR \mathcal{A} that uses the lazy time sampling. The encoding is sound and faithful (see [11] for more details).

We then recall the result that PSPACE and co-PSPACE are the same complexity class. Thus the realizability problem is PSPACE-hard.

D Realizability PSPACE upper bound proof (Theorem 1)

Proof. Let \mathcal{A} be a timed MSR constructed over finite alphabet Σ with J predicate symbols and E constant and function symbols. Let \mathcal{CS} be a critical configuration specification constructed over Σ and \mathcal{S}_0 be a given initial configuration. Let m be the number of facts in the initial configuration \mathcal{S}_0 , k an upper bound on the size of facts, and D_{max} a natural number that is an upper bound on the numeric values appearing in \mathcal{S}_0 , \mathcal{A} and \mathcal{CS} .

We propose a non-deterministic algorithm that accepts whenever there is a compliant trace starting from \mathcal{S}_0 in which time tends to infinity and which uses the lazy time sampling. We then apply Savitch's Theorem to determinize this algorithm.

In order to obtain the PSPACE result we rely on the equivalence among configurations which enables us to search for traces over δ -configurations [Appendix A] instead of searching for traces over concrete configurations. Furthermore, we rely on the assumption that functions \mathcal{N} and \mathcal{X} run in PSPACE to return 1 when a rule $r \in \mathcal{A}$ is applicable to a given δ -configuration, and when a δ -configuration is critical with respect to \mathcal{CS} , respectively. Additionally, we assume that the lazy time sampling is specified as a function \mathcal{T} which runs in PSPACE. \mathcal{T} takes a δ -configuration and a timed MSR and returns 1 when the tick must be applied and returns 0 when it must not be applied according to the lazy time sampling.

Because of Lemma 2, in the search for compliant traces, it suffices to consider traces of size bounded by the number of different δ -configurations, $L_\Sigma(m, k, D_{max})$ (stored in binary).

Let i be a natural number such that $0 \leq i \leq L_\Sigma(m, k, D_{max}) + 1$. The algorithm starts with $i = 0$ and W_0 set to be the δ -configuration of \mathcal{S}_0 and iterates the following sequence of operations:

1. If W_i is representing a critical configuration, *i.e.*, if $\mathcal{X}(W_i) = 1$, then return FAIL, otherwise continue;
2. If $i > L_\Sigma(m, k, D_{max}) + 1$, then ACCEPT; else continue;
3. If $\mathcal{T}(W_i, \mathcal{A}) = 1$ then replace W_i by W_{i+1} obtained from W_i by applying the *Tick* rule; Otherwise guess non-deterministically an instantaneous action, r , from \mathcal{A}

applicable to W_i , *i.e.*, such an action r that $\mathcal{N}(r, W_i) = 1$. If so replace W_i with the δ -configuration W_{i+1} resulting from applying the action r to the δ -configuration W_i . Otherwise FAIL;

4. Set $i = i + 1$.

We now show that this algorithm runs in polynomial space. The greatest number reached by the counter is $L_\Sigma(m, k, D_{max})$, which stored in binary encoding takes space bounded by:

$$\log(L_\Sigma(m, k, D_{max}) + 1) \leq m \log(J) + (m - 1) \log(D_{max} + 2) + mk \log(E + 2mk).$$

Therefore, in order to store the values of the step-counter, one only needs space that is polynomial in the given inputs.

Also, any δ -configuration, W_i can be stored in space that is polynomial to the given inputs. Namely, since W_i is of the form $[Q_1, \delta_{Q_1, Q_2}, Q_2, \dots, Q_{m-1}, \delta_{Q_{m-1}, Q_m}, Q_m]$ and values of the truncated time differences, $\delta_{i,j}$, are bounded, W_i can be stored in space $mk + (m - 1)(D_{max} + 2)$ which is polynomially bounded with respect to the inputs.

Finally, in step 3. algorithm needs to store the action r . This is done by remembering two configurations, while moving from one δ -configuration to another is achieved by updating the facts, updating the positions of facts and the corresponding truncated time differences and continue. Hence, step 3. can be performed in space polynomial to $m, k, \log_2(D_{max})$ and the sizes of \mathcal{N} and \mathcal{T} .

E Survivability PSPACE upper bound proof (Theorem 2)

Proof. We adapt the proof of Theorem 1 to survivability problem using the same notation and making the same assumptions.

In order to prove that \mathcal{A} satisfies survivability with respect to the lazy time sampling, \mathcal{CS} and \mathcal{S}_0 , we need to show that all infinite traces \mathcal{P} starting from \mathcal{S}_0 are compliant with respect to \mathcal{CS} . Since \mathcal{A} is progressive, in any infinite trace time necessarily tends to infinity, as per Proposition 3.

Based on our bisimulation result [Appendix A] we propose the search algorithm over δ -configurations instead of concrete configurations. We rely on Lemma 2 and search only for traces of size bounded by the number of different δ -configurations, $L_\Sigma(m, k, D_{max})$.

In order to prove survivability we first check realizability by using the algorithm given in the proof of Theorem 1. Notice that this algorithm is in PSPACE with respect to the inputs of survivability as well.

Next we show that no critical configuration is reachable from \mathcal{S}_0 using the lazy time sampling. The following algorithm accepts when a critical configuration is reachable, and fails otherwise. It begins with $i = 0$ and W_0 set to be the δ -configuration of \mathcal{S}_0 and iterates the following sequence of operations:

1. If W_i is representing a critical configuration, *i.e.*, if $\mathcal{X}(W_i) = 1$, then return ACCEPT, otherwise continue;
2. If $i > L_\Sigma(m, k, D_{max})$, then FAIL; else continue;
3. If $\mathcal{T}(W_i, \mathcal{A}) = 1$ then replace W_i by W_{i+1} obtained from W_i by applying the *Tick* rule; Otherwise guess non-deterministically an instantaneous action, r , from \mathcal{A}

applicable to W_i , *i.e.*, such an action r that $\mathcal{N}(r, W_i) = 1$. If so replace W_i with the δ -configuration W_{i+1} resulting from applying the action r to the δ -configuration W_i . Otherwise continue;

4. Set $i = i + 1$.

We take advantage of the fact that PSPACE, NPSPACE and co-PSPACE are all the same complexity class [22] and determinize the above algorithm and then switch the ACCEPT and FAIL. The resulting algorithm returns ACCEPT if and only if no critical configuration is reachable from the given initial configuration using the lazy time sampling.

The proof that above algorithms run in polynomial space is very similar to that proof relating to Theorem 1.

F n -time-bounded realizability is in NP (Theorem 3)

Let \mathcal{A} be a timed MSR constructed over finite alphabet Σ with J predicate symbols and E constant and function symbols. Let \mathcal{CS} be a critical configuration specification constructed over Σ and \mathcal{S}_0 be a given initial configuration. Let m be the number of facts in the initial configuration \mathcal{S}_0 , k an upper bound on the size of facts, and D_{max} a natural number that is an upper bound on the numeric values appearing in \mathcal{S}_0 , \mathcal{A} and \mathcal{CS} .

Moreover, assume that the function \mathcal{N} , \mathcal{X} and \mathcal{T} run in polynomial time with respect to the size of \mathcal{S}_0 . We show that we check in polynomial time whether a given trace \mathcal{P} is compliant and has exactly n -ticks. Because of Lemma 3, we know that traces have size of at most $(n + 2) * m + n$. Recall n is fixed. Set $i := 0$ and $ticks := 0$. Let W_i be the configuration at position i in \mathcal{P} . Iterate the following sequence of instructions:

1. if $i > (n + 2) * m + n$ then FAIL;
2. if $\mathcal{X}(W_i) = 1$ then FAIL;
3. if $ticks$ is equal to n , then ACCEPT;
4. if $\mathcal{T}(W_i) = 1$, then apply the Tick rule to W_i obtaining the configuration W_{i+1} and increment both $ticks$ and i ;
5. otherwise if $\mathcal{T}(W_i) \neq 1$, then guess non-deterministically a rule r , such that $\mathcal{N}(r, W_i) = 1$, apply this rule r to W_i , obtaining W_{i+1} , and increment i .

Since the size of facts is bounded and the number of facts in any configuration of the trace is m , all steps are done in polynomial time.

G Encoding of 3-SAT

Assume we are given a formula $F = (l_{11} \vee l_{12} \vee l_{13}) \wedge \dots \wedge (l_{n1} \vee l_{n2} \vee l_{n3})$.

We construct an initial configuration \mathcal{S}_0 and a progressive timed MSR \mathcal{A} that checks whether F is satisfiable or not. For readability, in the rules below, we do not explicitly write the set of constraints \mathcal{C}_r as per Definition 3. This set is implicitly assumed. For each variable v_i in F , we include the rules in \mathcal{A} :

$$\begin{aligned} Time@T, V_i@T_i \mid T \geq T_i &\longrightarrow Time@T, A_i@(T + 1) \\ Time@T, V_i@T_i \mid T \geq T_i &\longrightarrow Time@T, B_i@(T + 1) \end{aligned}$$

These rules rewrite the fact V_i to the fact A_i denoting true, or to the fact B_i denoting false. Intuitively, these rules construct an interpretation for the variables in F .

Now, we include the following rules which take an interpretation and reduce F accordingly:

$$\begin{aligned} & Time@T, A_k@T_1, I_{(v_k \vee l_{j_2} \vee l_{j_3}) \wedge C}@T_2 \mid T \geq T_1, T \geq T_2 \longrightarrow Time@T, A_k@T_1, I_C@(T+1) \\ & Time@T, A_k@T_1, I_{(l_{j_1} \vee v_k \vee l_{j_3}) \wedge C}@T_2 \mid T \geq T_1, T \geq T_2 \longrightarrow Time@T, A_k@T_1, I_C@(T+1) \\ & Time@T, A_k@T_1, I_{(l_{j_1} \vee l_{j_2} \vee v_k) \wedge C}@T_2 \mid T \geq T_1, T \geq T_2 \longrightarrow Time@T, A_k@T_1, I_C@(T+1) \end{aligned}$$

By inspection, the constructed \mathcal{A} is progressive timed MSR. We also have a polynomial number of rules, namely, $(2 \times p + 6 \times n)$ rules and a total of $(3 \times p + n + 1)$ predicates, where p and n are respectively the number of variables and clauses in F .

The initial configuration is $\mathcal{S}_0 = \{Time@0, V_1@0, \dots, V_p@0, I_F@0, Start@0\}$.

The fact $Start$ is never rewritten and is used to specify the critical configuration specification as follows: $\langle \{Start@T_1, I_C@T_2\}, \{T_2 \geq T_1 + n\} \mid C \neq \top \rangle$, where \top is the empty clause formula.

It is easy to see that our encoding is sound and complete: a configuration with the fact I_\top will be reached if and only if F is satisfiable. Moreover, there is a trace with exactly n ticks if and only if F is satisfiable. Before the first tick, we set all variables as true or false. We advance time. In the following n ticks, we use the rules above to evaluate I_F , the interpretation of the formula F . If F is not satisfiable, then no trace with n ticks will be compliant. If F is satisfiable, then there is a trace with n ticks that is compliant.